

A Deep Learning–Based Model for User Purchase Behavior Prediction in E-Commerce Platforms

Jing Zhao

How to cite: Zhao J. A Deep Learning–Based Model for User Purchase Behavior Prediction in E-Commerce Platforms. Textile & Leather Review. 2026; 9:1958-1978.

<https://doi.org/10.31881/TLR.2026.1958>

How to link: <https://doi.org/10.31881/TLR.2026.1958>

Published: 25 April 2026



A Deep Learning–Based Model for User Purchase Behavior Prediction in E-Commerce Platforms

Jing Zhao

Department of Economic and Trade, Shijiazhuang College of Applied Technology, Shijiazhuang, Hebei, 050800, China
13780510842@163.com

Article

<https://doi.org/10.31881/TLR.2026.1958>

Published 25 April 2026

ABSTRACT

Accurately predicting user purchase behavior remains a fundamental challenge for e-commerce platforms, particularly in categories such as apparel and textiles where exploratory browsing is prevalent. Although prior studies have predominantly relied on purchase intentions or proxy indicators, such measures often inadequately capture realized purchasing outcomes. To address this limitation, this study proposes a deep learning–based framework for session-level purchase behavior prediction that jointly incorporates platform engagement signals and customer attributes. Using large-scale post-filtered, high-intent e-commerce session data characterized by class imbalance, a range of machine learning and deep learning models were systematically evaluated within a unified experimental framework. The empirical results indicate that platform engagement features make a stronger contribution to purchase prediction than customer attributes, while customer attributes provide complementary contextual information. Compared with traditional machine learning approaches, the proposed deep learning model achieved the strongest overall predictive performance under the examined experimental setting. The present study focuses on post-filtered high-intent sessions and should therefore be interpreted as purchase prediction within decision-relevant shopping traffic rather than full-traffic platform prediction. Overall, these findings underscore the effectiveness of jointly modeling short-term behavioral signals and customer-level contextual attributes summarizing historical tendencies in predicting purchase outcomes. From a practical standpoint, the proposed framework facilitates decision support for operational tasks that require accurate session-level purchase prediction in e-commerce platforms. This study thus provides empirical evidence supporting the value of deep learning for modeling complex user behavior and contributes to the advancement of data-driven purchase prediction in online retail environments.

KEYWORDS

user purchase behavior, deep learning, session-level prediction, customer attributes, textile e-commerce

INTRODUCTION

Online retail platforms have become a primary channel for consumer purchasing activity. This shift is particularly pronounced in the textile and garment sectors, where digital marketplaces have significantly reshaped traditional consumption patterns. A growing body of research reports that an increasing share of commercial transactions now occurs through digital marketplaces rather than traditional brick-and-mortar stores [1,2]. As e-commerce platforms evolve, their operational priorities have shifted beyond simply enabling online transactions toward improving the conversion of large volumes of browsing activity into completed purchases [3,4]. Unlike offline retail settings, where store visits often signal a relatively strong purchase intention, user behavior on e-commerce platforms is frequently characterized by extensive browsing and comparison—a pattern frequently observed in apparel shopping where users evaluate style, material, and fit—often without an eventual transaction [2,5]. These interactions, commonly recorded as clickstream or session-level behavioral data, form a rich source of information that reflects users' preferences, decision-making processes, and responses to platform design and stimuli [1,6]. As a result, predicting whether a purchase will occur within a given high-intent shopping session has become a central challenge for platform operators, given its direct relevance to customer targeting, demand forecasting, and value-oriented decision support [3,7].

Major e-commerce platforms increasingly rely on behavioral data to tailor landing pages, ranking algorithms, and promotional content with the objective of converting exploratory browsing into purchase actions, particularly within sessions that already exhibit elevated purchase intent [2,8]. Accordingly, advances in user-level and session-level purchase prediction not only contribute to a deeper understanding of online consumer behavior but also have clear operational and economic implications for contemporary e-commerce platforms, especially for those dealing with heterogeneous and seasonal product categories like textiles [2,9,10].

Prior research on purchase behavior in e-commerce environments can be broadly grouped into two main streams. The first stream concentrates on behavioral interactions between users and platforms, examining how browsing activity, session-level engagement, and navigation patterns influence purchase outcomes [1,2]. Within this line of work, clickstream data are widely used to capture detailed interaction sequences and to infer latent shopping intentions [1,3]. Existing studies show that behavioral traces extracted from online sessions can be used to distinguish different shopping strategies and to construct predictive models linking platform visits to conversion outcomes [2,3]. In addition, platform design choices and interface characteristics have been shown to affect user trust and decision-making, thereby indirectly shaping purchase behavior [7,9].

The second stream adopts a customer-centric perspective, emphasizing user-specific characteristics rather than platform interactions alone. Research in this area highlights the role of factors such as historical purchasing behavior, customer lifetime value, and accumulated engagement indicators in shaping online purchase decisions [2,5]. Incorporating customer-level attributes allows platforms to differentiate users according to purchasing propensity and economic value, supporting more effective personalization and targeted marketing strategies [7]. Beyond behavioral history, demographic and contextual factors have also been examined to explain heterogeneity in purchasing behavior, with evidence suggesting that user-specific affective tendencies and other individual differences can influence online shopping decisions and cart abandonment behavior [5,9,11]. More recent work further underscores the importance of identifying high-value customers through longitudinal behavior modeling combined with predictive analytics to support retention and revenue optimization [12,13].

Taken together, these research streams highlight the importance of customer-centric information in understanding purchase behavior, while also revealing the limitations of approaches that focus exclusively on either behavioral interactions or stable user attributes. In contrast to studies that rely primarily on intention-based proxies or restricted behavioral signals, the present study adopts a session-level perspective centered on realized purchase outcomes and examines the combined contribution of platform engagement signals and customer attributes to purchase behavior prediction. The study does not aim to distinguish casual browsing from genuine shopping intent at the full-traffic platform level; instead, it focuses on purchase prediction within a subset of sessions that already exhibit substantive pre-checkout engagement.

RELATED WORK

Role of Platform Engagement in Determining Purchase Behavior

User engagement with e-commerce platforms constitutes the earliest observable stage of the online purchase decision process, during which awareness and consideration gradually take shape through browsing and interaction prior to any transaction [1,2]. Rather than being limited to traditional notions of customer satisfaction or loyalty, platform engagement has increasingly been viewed as a strategic asset that shapes conversion behavior and long-term value creation, particularly within sessions that exhibit meaningful purchase intent [2,7]. From a managerial standpoint, engagement is ultimately meaningful only insofar as it leads to completed purchases, rather than intermediate behavioral indicators such as clicks or dwell time alone [3,13]. Accordingly, a substantial body of prior research treats platform engagement as a key antecedent of online

purchase behavior, particularly in digital settings where purchasing decisions are mediated by platform design, interaction mechanisms, and data-driven personalization strategies [1,4]. To organize this line of research, existing studies commonly examine platform engagement from two complementary perspectives: platform-side engagement attributes and customer-side attributes [1,6].

Platform Engagement Attributes

From the platform perspective, user engagement is shaped by a range of interaction mechanisms, including recommendation systems, search functionality, content presentation, promotional strategies, and session-level personalization [8,13]. These design and interaction features influence how users navigate the platform, evaluate available products, and progress through successive stages of the decision-making process within a session. As a result, behavioral traces such as clickstream sequences, add-to-cart actions, and responses to promotional events are widely used as observable signals for modeling engagement dynamics and predicting purchase outcomes [3,4,9]. More recent research suggests that platform-driven engagement has become increasingly dynamic in contexts such as livestream commerce and targeted advertising, where sequential interactions and contextual cues play a decisive role in shaping conversion behavior [4,13]. However, a parallel line of work has shown that aggregated session-level representations remain effective for predicting purchase outcomes when the analytical objective focuses on completion decisions rather than interaction sequences.

Customer Attributes

While platform-side attributes shape the immediate interaction environment, customer-side attributes introduce substantial heterogeneity in engagement patterns and purchasing behavior. Prior research highlights the role of accumulated customer characteristics—such as historical purchasing behavior, value segmentation, and preference stability—in explaining differences in purchase propensity across users [1,7,12]. Even under identical platform conditions, customers may respond differently to engagement stimuli, reflecting variations in experience, loyalty, and economic value. In this sense, customer attributes complement platform-side signals by capturing stable, user-specific factors that influence purchase decisions over time, underscoring the importance of considering both perspectives when modeling online purchase behavior [6,12].

Online Purchase Prediction

Predicting whether users ultimately complete a purchase has long been a central topic across disciplines such as marketing, information systems, and data-driven platform research, given its direct implications for revenue generation and managerial decision-making [2,3]. Existing studies approach this problem from diverse angles,

including analyses of user motivations, platform capabilities, and patterns of online activity that precede purchasing decisions [1,7,13]. Despite these contributions, many studies rely on proxy measures—such as purchase intention, click-through behavior, or cart activity—that do not always correspond to actual purchase outcomes in real-world settings [3,9]. This gap reflects a fundamental limitation, as stated intentions and intermediate actions may diverge from final purchasing behavior under realistic constraints. In response, a growing body of work has begun to utilize real-world data sources, including clickstream logs and transaction records, to examine purchase behavior at the session level and to capture genuine conversion events more directly [2,4,14].

Analytical Methods for Purchase Prediction

A wide range of analytical methods has been applied to the purchase prediction problem, including regression-based models, decision trees, ensemble learning approaches, and shallow neural networks [6,10,13]. These techniques are widely used due to their interpretability and relative ease of implementation, but they often struggle to capture the complex and nonlinear relationships present in high-dimensional behavioral data [6,9]. Recent advances in deep learning have therefore attracted increasing attention, as deep models are capable of learning hierarchical feature representations from large-scale, heterogeneous data and adapting to dynamic behavioral patterns [2]. Empirical studies show that deep learning approaches can outperform conventional machine learning methods in modeling user interactions and predicting purchase-related outcomes in complex e-commerce environments [4,8]. Nevertheless, existing research has yet to establish a unified predictive framework that systematically integrates heterogeneous engagement signals with customer attributes, which motivates the methodological focus of the present study.

DATA DESCRIPTION

A major challenge in applying big data analytics to purchase prediction lies in the limited availability of reliable transaction-level data that capture realized purchasing behavior. Consequently, much of the existing literature relies on proxy measures—such as purchase intention, click-through behavior, or cart activity—to approximate customers' purchase decisions on e-commerce platforms [2,3,9,10,14]. While such proxies provide useful behavioral cues, recent studies increasingly demonstrate that models grounded in actual purchase outcomes yield more robust and practically meaningful insights than those based solely on intention-based or intermediate indicators [3,1,4]. In line with this perspective, the dataset used in this study is constructed to model

realized purchase outcomes at the session level, consistent with recent work in e-commerce analytics and decision support systems [1,2,4].

The data were collected from an online e-commerce platform operating as a multi-vendor, general-purpose marketplace, and all records were anonymized prior to analysis. The dataset integrates historical purchase information with a set of variables commonly adopted in session-level purchase prediction, combining customer characteristics with detailed browsing and interaction behaviors [1,2,7]. To ensure analytical relevance, the analysis focuses on active, intent-driven shopping sessions, defined as sessions in which users exhibit substantive engagement with product offerings rather than purely informational or immediately abandoned visits. This definition is based solely on observable pre-checkout interaction patterns and does not rely on purchase outcomes or post-transaction information.

In total, the dataset comprises 412,768 intent-driven sessions generated by 102,347 registered customers. Purchases occurred in 106,494 sessions (25.8%), while the remaining 306,274 sessions (74.2%) did not result in a transaction. The observed class distribution reflects the session selection strategy and operational characteristics of the platform, and should therefore be interpreted in the context of decision-relevant user traffic rather than full-traffic browsing behavior. Accordingly, the reported purchase proportion represents the conditional probability of purchase within the analyzed session subset and does not correspond to the baseline conversion rate of a full-traffic e-commerce environment. Accordingly, the present study should be interpreted as purchase prediction within a filtered high-intent session subset rather than as a direct evaluation of model performance on full-traffic e-commerce sessions dominated by casual browsing.

Each observation corresponds to a single user session and records in-session activities such as browsing behavior, product clicks, cart-related interactions, and purchase outcomes. To avoid information leakage, all session-level features were constructed exclusively from user interactions observable prior to any irreversible purchase commitment, such as checkout confirmation or payment submission. Variables including cart size, cart value, and number of products clicked were computed solely from behaviors observable before checkout confirmation or payment submission, regardless of whether a purchase ultimately occurred. No post-purchase or outcome-dependent information was included in the feature set.

Purchase sessions on the platform include both conventional cart-based checkout workflows and direct-purchase mechanisms (e.g., “buy-now” actions), in which explicit add-to-cart interactions may not occur prior to transaction completion. As a result, cart-related variables may remain zero in some purchase sessions because

no explicit cart formation step is required. Therefore, a cart size or cart value of zero should not be interpreted uniformly as indicating weak purchase intent. In the present study, cart-related variables are treated as session-level behavioral descriptors rather than universally applicable indicators of purchase intention across all transaction pathways. Within the aggregated session representation used in this study, cart-related variables are retained as one subset of transaction-related behavioral descriptors, but they are not assumed to function as universally valid primary signals across all transaction pathways. An overview of platform engagement and purchase outcomes is provided in Table 1. Although the dataset contains detailed session-level behavioral information, the analysis adopts aggregated session representations rather than modeling fine-grained event sequences, treating each session as a unified decision context in line with common practice in session-level purchase prediction studies [2,3,4].

Table 1. Platform engagement and purchasing behavior of customers

Description	Frequency
Number of prospective unique customers	102,347
Number of sessions	412,768
Number of sessions with a purchase	106,494
Number of sessions without a purchase	306,274

Note: Statistics are computed over post-filtered high-intent shopping sessions.

The relatively high proportion of purchase sessions reflects the behavioral filtering strategy used to identify high-intent sessions, rather than overall platform traffic characteristics. The dataset focuses on active sessions generated by registered customers in a mature e-commerce environment, where a substantial share of sessions corresponds to purposeful shopping behavior rather than casual browsing. Therefore, the observed class distribution is consistent with the operational context of the platform rather than an artifact of data construction.

To further characterize the sample, Table 2 reports the demographic attributes of customers represented in the data, while Table 3 presents a detailed description of the variables used in this study along with their descriptive statistics. Overall, the dataset includes two broad categories of variables: platform engagement attributes, which primarily describe session-level interaction behavior, and customer attributes, which capture customer-level contextual characteristics such as demographics, historical purchasing patterns, and account-related information.

Table 2. Demographics of the data

Demographic traits	Statistic
Gender (Female)	71.8%
Gender (Male)	28.2%
Age – Mean	42.6
Age – Std. dev.	11.4
Age (Female) – Mean	41.9
Age (Female) – Std. dev.	11.1
Age (Male) – Mean	44.1
Age (Male) – Std. dev.	12.3

Note: Demographic statistics are reported for customers who generated post-filtered high-intent shopping sessions.

Table 3. Descriptive statistics of variables used in this research

Variables	Minimum	Maximum	Mean	Std. Dev.
Platform engagement attributes				
Time of session (hour)	0	23	14.2	4.6
Day of week of session	1	7	–	–
Number of log-ins of customer	1	7	2.35	0.84
Number of products clicked	1	118	11.9	14.8
Lowest price of product clicked	1.0	398.5	28.7	61.5
Highest price of product clicked	1.0	915.2	96.4	142.7
Cart size	0	16	1.24	2.10
Cart value	0	1,215.6	62.8	118.9
Customer attributes				
Customer account score	92	645	478.6	127.4
Customer account lifetime (days)	3	720	141.6	114.1
Number of past purchases	0	96	5.7	9.8
Age of customer	18	82	42.5	11.3
Gender of customer	0	1	–	–
Days elapsed since last purchase	1	760	71.6	92.4

METHODOLOGY

This study adopts a data-driven modeling framework to predict user purchase behavior at the session level on e-commerce platforms. Prior research indicates that online purchase outcomes are jointly influenced by short-term behavioral interactions observable within a session and longer-term customer characteristics accu-

mulated through historical activity, with effects that are often nonlinear and context dependent [1,2]. Rather than relying on intention-based proxies, the proposed framework focuses explicitly on realized purchase outcomes and leverages session-level behavioral data. To accommodate the heterogeneity and scale of such data, the methodology integrates systematic data preprocessing, standardized model training procedures, and a deep learning–based predictive architecture, consistent with recent advances in e-commerce analytics [3,4,5]. The overall workflow is designed to support reliable performance estimation and fair comparison across alternative modeling approaches.

An overview of the modeling process is shown in Figure 1.

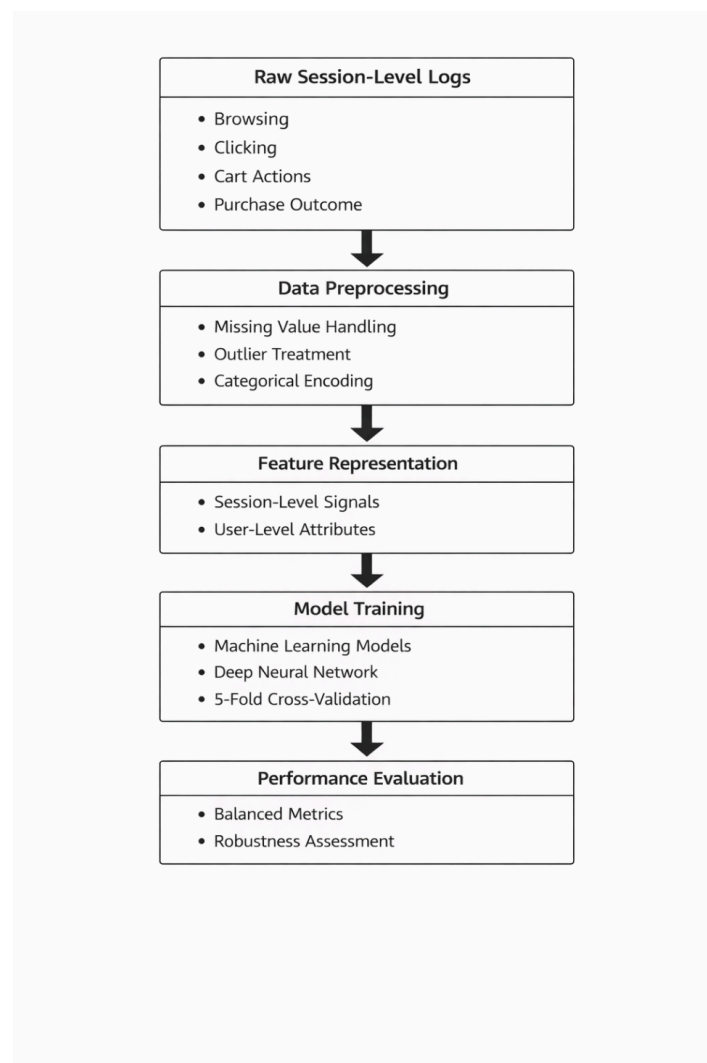


Figure 1. Overall modeling framework for session-level purchase behavior prediction

Data Preprocessing

Behavioral logs collected from e-commerce platforms typically contain noisy, sparse, and heterogeneous information, which may adversely affect predictive performance if not properly handled [1,7]. Accordingly, several preprocessing steps were applied prior to model construction to improve data quality and learning stability. Variables exhibiting extremely low variance or excessive missing values were removed, as such attributes contribute limited predictive information and may introduce unnecessary noise [5,9].

For the remaining variables, missing values were handled using training-fold-specific, model-aware imputation strategies that preserve relationships among observed variables, rather than simple distribution-based substitution. This approach allows the imputation process to exploit interdependencies among behavioral signals, which is particularly important in session-level purchase modeling [2,4]. Continuous features with heavy-tailed distributions were examined for extreme observations. Instead of discarding entire sessions, outliers were adjusted using mild clipping strategies to reduce numerical instability while retaining informative behavioral patterns [5].

All features were carefully time-aligned to ensure that only interactions preceding the purchase decision were included. In particular, session-level features were computed using logs recorded strictly before any irreversible commitment actions, such as checkout confirmation or payment submission. This design prevents information leakage and ensures that model predictions reflect genuine decision uncertainty rather than post-outcome signals. All preprocessing operations were performed independently within each cross-validation fold to avoid unintended information sharing across data partitions.

Categorical attributes were transformed into numerical representations suitable for both machine learning and deep learning models. Encoding strategies were selected to avoid introducing artificial ordinal relationships and to ensure consistency across different model classes [10]. After preprocessing, the final feature set consisted of a unified numerical representation of session-level engagement signals and customer-level contextual attributes, which served as input to all predictive models.

Model Training and Analysis

Model training was conducted under a unified experimental protocol to ensure robustness and comparability of results. Both machine learning and deep learning models require the specification of hyperparameters that substantially influence convergence behavior and predictive performance [2,10]. Reliance on a single train-test split may lead to unstable or overly optimistic performance estimates, particularly in large-scale

behavioral datasets [3]. To mitigate this risk, a five-fold cross-validation scheme was employed throughout the analysis [14].

Within each fold, hyperparameters were tuned using a grid-based search strategy based on validation performance within the training data, thereby reducing sensitivity to random data partitions and limiting the risk of hyperparameter overfitting [5]. Feature selection and refinement were conducted as part of the training process to reduce redundancy and enhance learning stability in high-dimensional behavioral data. Rather than emphasizing explicit feature importance ranking, this study focuses on constructing compact and informative feature representations that support effective learning across different model classes.

The close alignment observed between training and validation performance across folds suggests that the adopted training procedure achieves satisfactory generalization under the considered experimental setting.

Deep Learning Architecture

To capture complex interactions between platform engagement signals and customer attributes, a feed-forward deep neural network was employed. This architectural choice is appropriate because the input features are aggregated at the session level rather than represented as sequential event streams, rendering recurrent or convolutional structures less directly applicable to the present modeling objective [2,6]. Deep neural networks have demonstrated strong capability in learning hierarchical feature representations from high-dimensional behavioral data in e-commerce contexts [3,10].

The network architecture consists of multiple fully connected layers, where earlier layers capture basic feature interactions and deeper layers model higher-order dependencies. Model parameters were optimized using stochastic gradient-based learning, which provides an efficient and scalable optimization strategy for large datasets [6,12]. Different network depths and widths were explored to balance expressive capacity against computational efficiency, following empirical practices reported in prior purchase prediction studies [2,3].

Customer attributes in this framework primarily serve as contextual priors summarizing longer-term engagement tendencies, rather than as fine-grained longitudinal behavioral sequences. This design allows the model to modulate session-level engagement effects using accumulated user context while avoiding strong assumptions about sequential user behavior.

To enhance generalization performance, regularization techniques such as dropout and early stopping were incorporated during training [6,10]. The final model adopts three fully connected hidden layers with 256, 128, and 64 neurons, respectively, using ReLU activation functions. Model training was conducted using the Adam

optimizer with a learning rate of 0.001 and a batch size of 256. The maximum number of training epochs was set to 100, with early stopping based on validation loss. Dropout with a rate of 0.3 was applied to mitigate overfitting. All models were implemented using standard deep learning libraries, and random seeds were fixed across experimental runs to ensure reproducibility. To account for class imbalance in purchase outcomes, class weights inversely proportional to class frequencies were applied during training.

Performance Metrics

Model performance was evaluated using multiple complementary metrics to provide a balanced assessment of purchase prediction effectiveness. Because purchase events remain moderately imbalanced even at the session level, accuracy alone may be misleading [9,14]. Accordingly, precision, recall, F1-score, and the area under the receiver operating characteristic curve (ROC-AUC) were reported to capture different aspects of classification performance. In addition, the Matthews Correlation Coefficient (MCC) was adopted due to its robustness to class imbalance and its ability to jointly account for all outcomes in the confusion matrix [5,9]. All evaluation metrics were averaged across the five cross-validation folds to ensure stable and reliable performance estimation.

RESULTS

Predictive Contribution of Engagement and Customer Attributes

To examine the relative predictive contribution of different feature groups, predictive models were first evaluated using platform engagement attributes and customer attributes separately. Across all evaluated models, platform engagement features consistently exhibited stronger predictive power than customer attributes alone. Session-level behavioral signals—such as product exploration intensity and aggregated transaction-related interaction patterns—provided informative cues for purchase outcomes beyond static demographic or account-level variables. However, cart-related variables should be interpreted with caution, because some purchase sessions may follow direct-purchase workflows without explicit add-to-cart actions. Accordingly, these variables should be understood as pathway-dependent components of the aggregated session representation rather than as universally applicable primary signals of purchase intent.

Models trained on the combined feature set achieved the highest overall predictive performance, indicating that engagement behaviors and customer attributes capture complementary aspects of the purchase decision process. Engagement variables primarily reflect short-term behavioral dynamics within individual sessions,

whereas customer attributes contribute contextual information derived from longer-term historical behavior patterns. These findings are consistent with prior studies highlighting the important role of behavioral data in online purchase prediction while also recognizing the supplementary value of user-level characteristics [1,2].

Model Robustness and Stability Across Validation Folds

Model robustness was assessed by examining performance variability across the five cross-validation folds. The deep learning model exhibits lower variance in both accuracy and F1-score compared to traditional machine learning approaches, suggesting relatively more stable generalization across different data partitions. In contrast, tree-based models show greater sensitivity to data splits, particularly in folds characterized by uneven purchase distributions.

Robustness is an important requirement for real-world e-commerce applications, where user behavior may vary substantially across time periods and customer cohorts. The observed stability of the deep learning model suggests a stronger ability to maintain consistent performance under heterogeneous behavioral patterns, which aligns with prior evidence that deep architectures are well suited for modeling high-dimensional interaction data [3].

Table 4. Variability of model performance across cross-validation folds

Model	Std. Dev. (Accuracy)	Std. Dev. (F1-score)
Decision Tree (DT)	0.034	0.031
Random Forest (RF)	0.029	0.027
Artificial Neural Network (ANN)	0.021	0.020
Deep Neural Network (DNN)	0.016	0.015

Comparative Performance of Predictive Models

The overall comparative performance of all predictive models is summarized in Table 5. All evaluated approaches demonstrate acceptable discriminative capability, as indicated by ROC-AUC values exceeding the random baseline. More pronounced differences emerge, however, when balanced performance metrics are considered.

Among the compared models, the deep learning framework achieved the strongest overall performance in terms of F1-score and Matthews Correlation Coefficient (MCC), indicating a more favorable trade-off between precision and recall under the current experimental setting. Tree-based models achieve reasonable recall, indicating sensitivity to purchase events, but this is accompanied by comparatively lower precision and

higher false positive rates. By contrast, the deep learning model achieves a more balanced trade-off between correctly identifying purchase events and limiting false alarms, a property that is particularly desirable for session-level purchase prediction in operational e-commerce settings [5].

Table 5. Overall performance comparison of predictive models

Model	Accuracy	Precision	Recall	F1-score	MCC	ROC-AUC
Decision Tree (DT)	0.84	0.68	0.72	0.70	0.59	0.79
Random Forest (RF)	0.86	0.71	0.75	0.73	0.63	0.82
Support Vector Machine (SVM)	0.83	0.66	0.70	0.68	0.56	0.78
Artificial Neural Network (ANN)	0.88	0.75	0.78	0.76	0.68	0.84
Deep Neural Network (DNN)	0.89	0.78	0.80	0.79	0.72	0.86

Note:

- Precision, Recall, and F1-score are reported for the purchase (positive) class.
- All metrics are computed under a consistent decision threshold and averaged across five cross-validation folds.
- MCC is reported due to its robustness under class imbalance.

Statistical Validation of Performance Differences

To examine whether performance differences among models are statistically meaningful, non-parametric tests were conducted on fold-level evaluation metrics obtained from the five-fold cross-validation procedure. Given the relatively small number of cross-validation folds, effect sizes are reported descriptively rather than interpreted as definitive measures of practical superiority. The analysis suggests differences in performance across models; however, given the limited number of cross-validation folds, these results should be interpreted as indicative rather than conclusive.

However, given the limited number of folds, these statistical results should be interpreted with caution. Rather than serving as definitive proof of superiority, the statistical analysis complements the observed consistency in performance metrics and variance reduction across folds. Taken together, the results provide converging evidence that the deep learning model achieved the strongest overall performance among the compared models under the examined experimental setting.

Summary of Key Findings

In summary, the empirical results show that the integration of session-level engagement signals and customer attributes supports effective purchase prediction under the examined setting. The evaluated deep learning framework achieved the strongest overall predictive performance among the compared models. These findings support the value of combining heterogeneous session-level behavioral information with user-level contextual attributes in e-commerce purchase prediction.

DISCUSSION

This study investigates session-level purchase behavior prediction in e-commerce platforms by jointly modeling platform engagement signals and customer attributes within a deep learning framework. Building on the empirical results presented in Section 5, this discussion interprets the findings from the perspectives of predictive mechanisms, model robustness, and practical relevance, and situates them within the broader literature on online purchase prediction.

Role of Engagement and User Attributes in Purchase Prediction

The results reported in Section 5.1 suggest that session-level engagement signals contribute more directly to purchase prediction than static customer attributes. This pattern is understandable because engagement variables capture immediate within-session behavioral dynamics, whereas customer attributes mainly provide broader contextual information accumulated over time. Session-level signals—such as product exploration intensity and transaction-related interaction patterns—provided useful behavioral cues associated with short-term purchasing behavior. However, cart-related variables do not carry a uniform behavioral meaning across all sessions, because some transactions may be completed through direct-purchase mechanisms without explicit add-to-cart actions. Accordingly, these variables should be understood as pathway-dependent components of the aggregated session representation rather than as universally applicable primary signals of purchase intent. This interpretation is consistent with prior studies emphasizing the predictive value of behavioral data in online retail environments [1,2].

At the same time, the performance gains achieved by combining engagement features with customer attributes suggest that historical and contextual user information remains an important complement. Customer attributes reflect accumulated contextual information about users' historical interactions with the platform that cannot be inferred from a single session alone. Taken together, these findings highlight the importance

of jointly modeling short-term behavioral signals and longer-term user context in session-level purchase prediction.

Robustness and Generalization of Predictive Models

Robustness across different data partitions is essential for real-world deployment in dynamic e-commerce settings. The deep learning model exhibits consistently lower performance variance across cross-validation folds than traditional machine learning approaches, indicating more stable generalization under varying data conditions. This stability suggests a stronger capacity to handle heterogeneity in user behavior and fluctuations in purchase distributions.

By contrast, tree-based models demonstrate greater sensitivity to data splits, which may limit their reliability when user behavior patterns evolve over time. The robustness advantage observed for the deep learning model can be attributed to its ability to model nonlinear interactions among heterogeneous features and to maintain stable performance under noisy behavioral signals [3].

Balanced Performance and Implications for Platform Decision-Making

The comparative results summarized in Table 5 show that, although all evaluated models achieve reasonable discriminative performance, the deep learning model performs better than alternative approaches on evaluation metrics such as the F1-score and the Matthews Correlation Coefficient. These results indicate a more favorable trade-off between correctly identifying purchase events and limiting false positive predictions.

From an operational perspective, controlling false positives is particularly important for e-commerce platforms, as inaccurate purchase predictions may lead to inefficient targeting strategies and suboptimal allocation of marketing resources. The observed performance of the deep learning model therefore supports its potential utility for operational decision support tasks that require a careful trade-off between precision and recall, such as ranking or targeting under similar platform conditions [5].

Methodological Implications and Limitations

Beyond predictive performance, the findings carry broader methodological implications. Under the specific platform conditions examined in this study, the results suggest that deep learning methods are applicable to session-level purchase prediction in settings characterized by class imbalance and feature heterogeneity [4]. These observations are based on session-level purchase prediction under the specific data characteristics examined in this study. An additional limitation is that the comparative analysis did not include stronger gradient-boosting baselines commonly used for tabular prediction tasks, such as XGBoost or LightGBM.

Therefore, the present findings should be interpreted as evidence within the examined model set rather than as definitive proof of the universal superiority of deep learning for tabular purchase prediction. An additional limitation concerns the interpretation of cart-related variables across heterogeneous transaction mechanisms. Specifically, sessions completed through direct-purchase workflows may retain zero cart size or cart value despite leading to a completed purchase. As a result, such variables cannot be assumed to carry a uniform behavioral meaning across all sessions. Future research may address this issue by explicitly modeling transaction pathway types when suitable labels are available. However, these methodological benefits are accompanied by increased computational demands.

Practitioners must therefore balance predictive performance against computational and deployment costs. Future research may explore lightweight architectures or hybrid modeling strategies that preserve predictive accuracy while improving scalability and resource efficiency.

IMPLICATIONS

Academic Implications

This study makes several contributions to the literature on e-commerce analytics and purchase behavior prediction. First, it advances prior research that largely relies on purchase intentions or proxy indicators by focusing on realized, session-level purchase outcomes. By modeling actual purchasing behavior rather than intention-based measures, the findings provide stronger empirical support for understanding how online engagement is associated with realized purchase outcomes [1,2].

Second, the study contributes methodologically through a systematic comparison of deep learning and traditional machine learning approaches within a unified experimental framework. The findings provide empirical support for the applicability of deep learning methods to session-level purchase prediction involving heterogeneous behavioral and contextual features [3]. Third, this study extends the application of deep learning to session-level purchase prediction in e-commerce settings by jointly modeling short-term engagement behaviors and customer-level contextual attributes summarizing historical tendencies. The observed contribution of engagement signals, together with the complementary role of customer attributes, opens avenues for future research on how immediate decision cues and accumulated user context interact to shape purchasing dynamics over time [5].

Managerial Implications

From a managerial perspective, the findings offer actionable insights for platform operators. The strong predictive value of session-level engagement signals suggests that session-level behavioral indicators should be prioritized when designing analytical tools to support recommendation or promotion-related decision-making. In competitive environments, data-informed actions before checkout completion may be more effective than strategies aimed solely at prolonging user engagement. At the same time, the complementary contribution of customer attributes underscores the importance of differentiating strategies across customer segments when historical contextual information is available. While short-term behaviors capture immediate intent, historical user information can support more differentiated analytical strategies across customer segments. Finally, the results highlight the need for selective data governance. Rather than indiscriminately accumulating data, firms should focus on high-value behavioral and contextual features that directly enhance purchase prediction, balancing predictive gains with privacy and computational considerations [4].

CONCLUSION

This study examined session-level user purchase behavior prediction on e-commerce platforms using a deep learning-based framework that integrates platform engagement signals with customer attributes. By emphasizing realized purchase outcomes rather than intention-based proxies, the proposed approach provides a more realistic and operationally relevant perspective on online purchasing behavior.

The empirical results showed that session-level engagement signals and customer attributes played complementary roles in purchase prediction, with engagement-related information contributing more directly under the examined experimental setting. Relative to the compared machine learning baselines included in this study, the deep learning framework achieved the strongest overall predictive performance under the examined experimental setting.

Beyond predictive performance, the findings contribute to both methodological understanding and practical application. Methodologically, the study provides empirical support for the value of jointly modeling heterogeneous session-level engagement signals and customer-level contextual attributes in e-commerce purchase prediction. Practically, the results suggest that such signals can support analytical decision-making for recommendation-, promotion-, and resource allocation-related tasks. In the garment retail context, this type of decision support may help mitigate high return rates and improve seasonal inventory planning through more accurate identification of purchase propensity.

Several limitations should be acknowledged. The analysis relies on historical session data from a specific e-commerce platform, which may constrain the generalizability of the findings. Future research could extend this work by incorporating domain-specific product features, such as textile material and style tags, temporal dynamics across sessions, cross-platform behavioral data, or alternative lightweight deep learning architectures to improve scalability and generalization.

Overall, this study advances data-driven understanding of session-level purchase prediction and highlights the potential practical value of deep learning for decision support in comparable online retail contexts characterized by intent-driven session data.

Availability of Data and Materials

The datasets used and/or analysed during the current study were available from the corresponding author on reasonable request.

Author Contributions

Jing Zhao designed, collected and analyzed the data, and drafted the manuscript. Jing Zhao conducted the study, critically revised the manuscript for important intellectual content, and gave final approval of the version to be published. Jing Zhao participated fully in the work, take public responsibility for appropriate portions of the content, and agreed to be accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved.

Conflicts of Interest

The author declares no conflict of interest.

Funding

This research received no external funding.

Ethics Approval and Consent to Participate

This survey was conducted in compliance with Ethics Committee of Shijiazhuang College of Applied Technology. Participants were informed of the study's purpose and data usage prior to participation, and responses were collected anonymously. No personally identifiable information was stored.

Acknowledgment

Not applicable.

REFERENCES

- [1] Kim S, Shin W, Kim HW. Predicting online customer purchase: The integration of customer characteristics and browsing patterns. *Decision Support Systems*. 2024; 177:114105. doi: 10.1016/j.dss.2023.114105
- [2] Chaudhuri N, Gupta G, Vallurupalli VV, Bose I. On the platform but will they buy? Predicting customers' purchase behavior using deep learning. *Decision Support Systems*. 2021; 149:113622. doi: 10.1016/j.dss.2021.113622
- [3] Wei S, Yang Z, Zhang J, Zeng Y, Li Q, Xiao Y. Click-through conversion rate prediction model of book e-commerce platform based on feature combination and representation. *Expert Systems with Applications*. 2024; 238:122276. doi: 10.1016/j.eswa.2023.122276
- [4] Lin Q, Jia N, Chen L, Zhong S, Yang Y, Gao T. A two-stage prediction model based on behavior mining in livestream e-commerce. *Decision Support Systems*. 2023; 169:114013. doi: 10.1016/j.dss.2023.114013
- [5] Xiao F, Chen SX, Chen ZY, Wang YN, Wang JQ. An interpretable data-driven approach for customer purchase prediction using cost-sensitive learning. *Engineering Applications of Artificial Intelligence*. 2024; 138:109344. doi: 10.1016/j.engappai.2024.109344
- [6] Yu B, Zhang J, Fu Y, Xu Z. A weighted heterogeneous graph attention network method for purchase prediction of potential consumers with multibehaviors. *Information Processing & Management*. 2025; 62(5):104175. doi: 10.1016/j.ipm.2025.104175
- [7] Sakalauskas V, Kriksciuniene D. Personalized advertising in e-commerce: Using clickstream data to target high-value customers. *Algorithms*. 2024; 17(1):27. doi: 10.3390/a17010027
- [8] Celik E, İlhan Omurca S. Skip-gram and transformer model for session-based recommendation. *Applied Sciences*. 2024; 14(14):6353. doi: 10.3390/app14146353
- [9] Waldmann F, Nápoles G, Salgueiro Y. Prediction of cart abandonment using imbalanced clickstream data: An online personalization approach. *Electronic Commerce Research*. 2025. doi: 10.1007/s10660-025-10021-3
- [10] Alotaibi S, Alotaibi B. A review of click-through rate prediction using deep learning. *Electronics*. 2025; 14(18):3734. doi: 10.3390/electronics14183734
- [11] Silva P, Fortes N, Bairrada C, Coelho F. Hit the Road, Cart: A Study on Affect and Online Shopping Cart Abandonment. *International Journal of Consumer Studies*. 2025; 49(3):e70060. doi: 10.1111/ijcs.70060
- [12] Kasemrat R, Kraiwanit T. Attention-enhanced LSTM for high-value customer behavior prediction: Insights from Thailand's e-commerce sector. *Intelligent Systems with Applications*. 2025; 26:200523. doi: 10.1016/j.iswa.2025.200523

- [13] Chan Z. Capturing conversion rate fluctuation during sales promotions: A novel historical data reuse approach. In Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23). 2023. doi: 10.1145/3580305.3599788
- [14] Chopra IP, Jebarajkirthy C, Jain T, Maseeh HI. Electronic shopping cart abandonment: What do we know and where should we be heading? *Electronic Markets*. 2024; 34:Article 25. doi: 10.1007/s12525-024-00697-0